Rule WLM018:   Multiple periods specified for server service class

Finding: CPExpert noticed that more than one period had been specified for a server service class.

Impact: This finding has a LOW impact on the performance of your system. The impact mostly results from increased (and unnecessary) overhead.

Logic flow: This a basic finding. There are no predecessor rules.

Discussion: A service class may be broken into multiple service class **periods**. Each service class has Period 1 automatically defined. Optionally, installations can define up to seven additional service class periods (although typically no more than two or three additional service class periods are defined). Each service class period can have its own performance goals, defined to the Workload Manager via the Workload Manager ISPF panels.

An address space (TSO transaction, batch job step, etc.) begins in Period 1 of the service class to which it is assigned. The address space transitions from Period 1 to Period 2 (and to subsequent periods), based upon the accumulation of "service" by the transaction. The "service" required by the address space is a combination of CPU resources, I/O resources, and memory resources.

The normal purpose of defining multiple service class periods is to give higher importance to interactive transactions, short batch job steps, etc. Overall response is decreased (and overall throughput is increased) when address spaces requiring relatively few resources do not compete for system resources with those address spaces requiring substantial resources[1].

If subsystems are installed which support Workload Manager reporting (e.g., CICS/ESA Version 4.1 or IMS/ESA Version 5), installations can define service classes which describe particular transaction types and specify performance goals for the transactions in the service class. All transactions entering the system which fall into the workload category described by the service class are associated with the service class.

---

[1]Note that with MVS/ESA SP5.1, the redesign of the MVS Dispatcher algorithms eliminated much of the potential conflict for CPU cycles (although other considerations apply). Please refer to Rule WLM025 for additional discussion in this area.

Rule WLM018.1

Please refer to Section 4 for discussion of the relationship between server service classes and transaction service classes.

Of particular importance to this finding is the fact that the Workload Manager normally uses the performance goal for the server service class only during address space start-up (and may use it during address space shutdown). During other times, the Workload Manager **uses the performance goals (and goal importance) of the transaction service classes being served**, and ignores the performance goal specified for the server.

Since the only purpose of having multiple periods for a service class is to specify different performance goals and/or goal importance for different categories of work, and since the Workload Manager ignores the performance goal and goal importance of the server, multiple periods should not be specified for server service classes. The changed performance goal or goal importance will be ignored by the Workload Manager, and the server will not "migrate" to lower service class periods. However, the system will incur additional overhead caused by the unused service class period(s). As examples of the overhead:

• Additional SRM control blocks are created and processed.

• Additional Workload Manager control blocks are created and processed.

• RMF requires additional processing of the control blocks.

• SMF Type 72(Subtype 3) records are written for each service class period defined (regardless of whether the service class periods are used by the Workload Manager).

CPExpert produces Rule WLM018 when it detects that multiple periods are defined for a server service class.

The following example illustrates the output from Rule WLM018:

```
 RULE WLM018:  MULTIPLE PERIODS WERE SPECIFIED FOR SERVER

    CPExpert noticed that the CICSRGN Service Class was defined to include
    multiple periods.  The CICSRGN Service Class is a server (e.g., a CICS
    or IMS region), and multiple periods cannot be used for servers.
    CPExpert suggests that you eliminate the multiple-period definition
    for the CICSRGN Service Class, since the unused periods will generate
    unnecessary overhead.
```

Suggestion: CPExpert suggests that you eliminate the multiple periods for the server service class. These periods are not used, but they cause unnecessary overhead.

Reference: MVS Planning: Workload Management

MVS/ESA(SP 5):	Chapter 8.2: Using Performance Periods
OS/390 (V1R1):	Chapter 8.2: Using Performance Periods
OS/390 (V1R2):	Chapter 8.2: Using Performance Periods
OS/390 (V1R3):	Chapter 8.2: Using Performance Periods
OS/390 (V2R4):	Chapter 8.3: Using Performance Periods
OS/390 (V2R5):	Chapter 8.3: Using Performance Periods
OS/390 (V2R6):	Chapter 8.3: Using Performance Periods
OS/390 (V2R7):	Chapter 8.3: Using Performance Periods
OS/390 (V2R8):	Chapter 8.3: Using Performance Periods
OS/390 (V2R9):	Chapter 8.3: Using Performance Periods
OS/390 (V2R10):	Chapter 8.3: Using Performance Periods
z/OS (V1R1):	Chapter 8.3: Using Performance Periods
z/OS (V1R2):	Chapter 8.3: Using Performance Periods
z/OS (V1R3):	Chapter 8.3: Using Performance Periods
z/OS (V1R4):	Chapter 8.3: Using Performance Periods